



#12 / Appeal Brief
T. R. K. 6/3/2004
6/3/2004

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

In re Application of:
David A. Wood et al.

Serial No. 09/487,529

Filed: January 19, 2000

For: System and Method for Enhancing
Communication Between Devices
in a Computer System

§ Group Art Unit: 2126
§
§ Examiner: Hoang, Phuong N.
§
§ Atty. Dkt. No.: 5181-38400
§ P4359

CERTIFICATE OF MAILING 37 C.F.R. § 1.8	
I hereby certify that this correspondence is being deposited with the U.S. Postal Service with sufficient postage as First Class Mail in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date indicated below:	
<u>B. Noel K. v. L.</u> Name of Registered Representative	
5-10-04	Signature
Date	Signature

RECEIVED

APPEAL BRIEF

MAY 17 2004

Technology Center 2100

Box AF
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir/Madam:

Further to the Notice of Appeal filed March 9, 2004, Appellants present this Appeal Brief. Appellants respectfully request that this appeal be considered by the Board

06/07/2004 TMCBETH 00000001 501505 09487529
of Patent Appeals and Interferences.

01 FC:1402 330.00 DA

I. REAL PARTY IN INTEREST

The subject application is owned by Sun Microsystems, Inc., a corporation organized and existing under and by virtue of the laws of the State of Delaware, and having its principal place of business at 4150 Network Circle, Santa Clara, CA 95054, as evidenced by the assignment recorded at Reel 010522, Frame 0250.

II. RELATED APPEALS AND INTERFERENCES

No other appeals or interferences are known which would directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 1 – 20 are pending in the present application and are the subject of this appeal. Claims 1 – 20 stand finally rejected under 35 U.S.C. § 102(e). A copy of Claims 1 – 20, as on appeal (incorporating all amendments), is included in the Appendix hereto.

IV. STATUS OF AMENDMEMNTS

No amendment to the claims has been filed subsequent to the final rejection. The Appendix hereto reflects the current state of the claims.

V. SUMMARY OF THE INVENTION

A sending device and a target device may be configured to communicate according to a packet-based communication protocol via a communications medium. For example, the sending device may be configured to convey a request to the target device. The target device may be configured to convey a response after receiving or processing the request from the sending device. The response conveyed from the target device may comprise an acknowledgment (ACK) or a negative acknowledgment (NAK) according to the communications protocol employed by the devices. *See page 4, lines 10-24.*

In certain situations, the target device may be temporarily unable to process the request from the sending device. These may be referred to as a “temporarily unavailable conditions” and may occur when the target device is handling another operation that temporarily prevents the processing of the request from the sending device. Such operations may include a temporarily loss of system resources (e.g., a dynamic reconfiguration of a node), a temporary lack of processing resources on the target device, or a lack of a valid virtual to physical address translation in cases where the contents of the request are to be written in the virtual address space of the target device’s node. In response to detecting a temporarily unavailable condition, target device may be configured to convey a negative acknowledgment (NAK) to the sending device. In one embodiment, the target device may be configured to convey different types of NAKs depending on the type of temporarily unavailable condition detected. *See page 4, line 26 – page 5, line 9.*

In one embodiment, a NAK may include a delay value that may be used by the sending device as a hint for determining how long to delay the resending of its request. Using the delay value, the sending device may advantageously resend its request at a time when the target device may be able to process the request, i.e., after sufficient time to allow the temporarily unavailable condition to be cleared at the target device. In certain configurations or for certain types of temporarily unavailable conditions, the sending device may be configured to ignore the delay value and independently determine when to resend its request. *See page 5, lines 10-16.*

In one embodiment, the target device may be configured to generate a delay value according to the type of operation that is causing a temporarily unavailable condition. In this manner, different delay values can be generated for different types of operations, because the amount of time necessary for the target device to clear the temporarily unavailable condition may vary among the different types of operations. Also, the target device may generate delay values according to a set value for each type of operation, a

programmed value for each type of operation, or a dynamically calculated value for each type of operation. The target device may be configured to store historical data from previous temporarily unavailable conditions and may calculate delay values from this data. The target device may also keep track of the number of outstanding responses it has sent for a particular temporarily unavailable condition. In doing so, the target device may convey delay values that indicate increasingly longer delay periods as the number of outstanding responses increases. Additionally, the delay value may be encoded to minimize the size and/or number of packets needed for the NAK. *See page 5, line 18 – page 6, line 1.*

In one embodiment, a policy layer may determine a retry limit for a particular request sent by the sending device. If the sending device resends the request in excess of the retry limit, the policy layer may be configured to detect an error and may initiate an error recovery mechanism based on the type of NAK most recently received from the target device. The type of NAK may allow for different error recovery mechanisms based on different types of temporarily unavailable conditions at the target device. In other embodiments, the policy layer may be configured to detect an error and may initiate an error recovery mechanism based on the delay value corresponding to the most recently received NAK from the target device. *See page 6, lines 15-23.*

VI. ISSUES

1. Whether claims 1-20 are anticipated by Chambers et al. (U.S. Patent 5,884,052) under 35 U.S.C. § 102(e).

VII. GROUPING OF CLAIMS

Claims 1, 2-3, 8-11, 15, and 16 stand or fall together.

Claim 19 stands or falls alone.

Claim 4 stands or falls alone.

Claims 5, 12, and 20 stand or fall together.

Claims 6, 13, and 18 stand or fall together.

Claims 7, 14, and 17 stand or fall together.

The above claim groupings are for purposes of this appeal only. The reasons why each group of claims is believed to be separately patentable are explained below in the Argument.

VIII. ARGUMENT

A. Claims 1, 2-3, 8-11, 15, and 16

The Examiner rejected claims 1, 2-3, 8-11, 15, and 16 under 35 U.S.C. § 102(e) as being anticipated by Chambers et al. (U.S. Patent 5,884,052) (hereinafter “Chambers”). Appellants respectfully traverse these rejections in light of the following remarks.

Contrary to the Examiner’s assertion, Chambers fails to teach or suggest, “said second device is configured to convey a response to said first device corresponding to said first request, and wherein said response includes a delay value corresponding to said temporarily unavailable condition” as recited in claim 1. In addition, Chambers fails to teach or suggest, “generating a delay value corresponding to said temporarily unavailable condition; and conveying a response corresponding to said first request to from said second device to said first device, wherein said response includes said delay value” as recited in claim 15.

Chambers teaches “a data transaction between an initiator PCI agent and a target PCI agent. The initiator PCI agent is a relatively high performance PCI agent having a low latency. The target PCI agent is a relatively slow performing, high latency PCI agent.” (Chambers, column 5, lines 16-20) In Chambers, the initiator PCI agent initially accesses the target PCI agent, but the target issues a retry due to the target’s data access latency, which causes the initiator to terminate the data transaction before completion. (Chambers, column 5, lines 24-27)

Additionally, Chambers teaches that the target PCI agent has a known internal access latency, e.g., a latency of 16 PCI clock cycles. The known internal access latency is characterized as a “delay input” and it is used to eliminate retries during the latency period. (Chambers, column 5, lines 44-51) In Chambers, a retry delay register, which is coupled to initiator PCI agent, receives and stores the “delay input” (i.e., the known

latency of the target PCI agent) during a configuration access (e.g., system startup) to implement the variable retry strategy. (Chambers, column 6, lines 44-59) The variable retry strategy is characterized by n, 2, 2, 2..., where n is the known latency of the target PCI agent. (Chambers, column 6, lines 6-8)

More specifically, as shown in Figure 7 of Chambers, the retry delay register is first configured with the “delay input” or the known latency of the target PCI agent during a configuration access. Next, the initiator PCI agent attempts an initial access to the high latency target PCI agent. The target PCI agent then sends a retry. (Figure 7 and col. 8, lines 5 – 24)

In support of his contention that Chambers teaches that the target sends the initiator a response that includes a delay value corresponding to a temporarily unavailable condition, the Examiner states, in the Final Office Action dated January 15, 2004, and in the Advisory Action dated April 6, 2004, that this feature is taught in col. 5, lines 45-55 of Chambers. Applicant respectfully submits that col. 5, lines 45-55, of Chambers teaches that a target device has a known internal access latency, which is referred to as the “delay input.” However, Chambers fails to teach that the “delay input” is sent in the response (i.e. retry) from the target device to the requesting device. Instead, Chambers teaches that the “delay input” is loaded into the retry delay register during a configuration access (e.g., system start up). (Chambers, Column 6, Lines 44-59, Column 7, Lines 25-27, Column 7, Lines 45-46, and Column 8, Lines 5-16) Also, Figure 7 of Chambers illustrates that the retry delay register is configured (i.e., the delay input is loaded) before the transaction between the requesting device and the target device (block 701 and col. 8, lines 5 – 24). Therefore, Chambers fails to teach or suggest, “said response includes a delay value corresponding to said temporarily unavailable condition” as recited in claim 1.

In addition to the arguments given above, Chambers fails to teach that the target device generates the delay value according to the type of temporarily unavailable

condition. In fact, Chambers fails to teach that the target generates any delay values. In Chambers, the “delay input” is a known internal access latency of a particular target device. (Chambers, Column 5, Lines 44-51) Therefore, Chambers fails to teach or suggest, “generating a delay value corresponding to said temporarily unavailable condition; and conveying a response corresponding to said first request to from said second device to said first device, wherein said response includes said delay value” as recited in claim 15.

Anticipation requires the presence in a single prior art reference disclosure of each and every element of the claimed invention, arranged as in the claim. *Lindemann Maschinenfabrik GmbH v. American Hoist & Derrick Co.*, 221 USPQ 481, 485 (Fed. Cir. 1984). The identical invention must be shown in as complete detail as is contained in the claims. *Richardson v. Suzuki Motor Co.*, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). Chambers clearly does not anticipate Appellants’ claimed invention.

In accordance, claims 1 and 15 are believed to patentably distinguish over the cited reference. Claims 2-3 and 16 are dependent on claims 1 and 15, respectively, and are therefore believed to patentably distinguish over the cited reference for at least the reasons given above. Claim 8 recites features similar to those of claim 1 as discussed above and is thus also believed to patentably distinguish over the cited reference for at least the same reasons given above, along with their respective dependent claims 9-11. Since the rejection is not supported by the teaching of the cited reference, Appellants respectfully request reversal of the Examiner’s rejection of claims 1, 2-3, 8-11, 15 and 16.

B. Claim 19

The Examiner rejected claim 19 under 35 U.S.C. § 102(e) as being anticipated by Chambers. Appellants respectfully traverse these rejections in light of the following remarks.

The rejection of claim 19 is unsupported by the cited reference for at least the reasons given above in Argument A. Furthermore, contrary to the Examiner's assertion, Chambers fails to teach or suggest "determining a type of said temporarily unavailable condition; and generating said delay value according to said type of said temporarily unavailable condition" as recited in claim 19.

Chambers teaches that the target PCI agent has a known internal access latency, i.e., the "delay input." (Chambers, column 5, lines 44-51) Chambers fails to teach or suggest determining a type of temporarily unavailable condition because Chambers is limited to detecting the known internal access latency of the target device. Also, Chambers fails to teach or suggest that the target device generates the delay value according to the type of temporarily unavailable condition. For example, in some embodiments of Applicant's invention, different delay values can be generated for different types of operations, because the amount of time necessary for the target device to clear the temporarily unavailable condition may vary among the different types of operations. Also, the target device may generate delay values according to a set value for each type of operation, a programmed value for each type of operation, or a dynamically calculated value for each type of operation. The target device may be configured to store historical data from previous temporarily unavailable conditions and may calculate delay values from this data. The target device may also keep track of the number of outstanding responses it has sent for a particular temporarily unavailable condition. In doing so, the target device may convey delay values that indicate increasingly longer delay periods as the number of outstanding responses increases. (Applicant's Specification, Page 5, Lines 18-29)

In fact, Chambers fails to teach or suggest that the target generates any delay values. Instead, Chambers teaches that the retry delay register, which is coupled to initiator PCI agent, receives and stores the "delay input" (i.e., the known latency of the target PCI agent) during a configuration access (e.g., system startup). (Chambers, column 6, lines 44-59) In accordance, claim 19 is believed to patentably distinguish over the

cited reference. Since the rejection is not supported by the teaching of the cited reference, Appellants respectfully request reversal of the Examiner's rejection of claim 19.

C. **Claim 4**

The Examiner rejected claim 4 under 35 U.S.C. § 102(e) as being anticipated by Chambers. Appellants respectfully traverse these rejections in light of the following remarks.

The rejection of claim 4 is unsupported by the cited reference for at least the reasons given above in Argument A. Additionally, Chambers fails to teach or suggest “said delay value corresponds to a first value in response to said temporarily unavailable condition corresponding to a first type of condition and wherein said delay value corresponds to a second value in response to said temporarily unavailable condition corresponding to a second type of condition” as recited in claim 4.

Chambers teaches that the target PCI agent has a known internal access latency, i.e., the “delay input.” (Chambers, column 5, lines 44-51) Chambers fails to teach or suggest that the delay value may vary depending on the type of detected temporarily unavailable condition. For example, in some embodiments of Applicant's invention, different delay values can be generated for different types of operations, because the amount of time necessary for the target device to clear the temporarily unavailable condition may vary among the different types of operations. Also, the target device may generate delay values according to a set value for each type of operation, a programmed value for each type of operation, or a dynamically calculated value for each type of operation. The target device may be configured to store historical data from previous temporarily unavailable conditions and may calculate delay values from this data. The target device may also keep track of the number of outstanding responses it has sent for a particular temporarily unavailable condition. In doing so, the target device may convey

delay values that indicate increasingly longer delay periods as the number of outstanding responses increases. (Applicant's Specification, Page 5, Lines 18-29)

Applicant respectfully submits that Chambers is limited to detecting the known internal access latency of the target device. Chambers teaches that the retry delay register, which is coupled to initiator PCI agent, receives and stores the "delay input" (i.e., the known latency of the target PCI agent) during a configuration access (e.g., system startup). In accordance, claim 4 is believed to patentably distinguish over the cited reference. Since the rejection is not supported by the teaching of the cited reference, Appellants respectfully request reversal of the Examiner's rejection of claim 4.

D. Claims 5, 12, and 20

The Examiner rejected claims 5, 12, and 20 under 35 U.S.C. § 102(e) as being anticipated by Chambers. Appellants respectfully traverse these rejections in light of the following remarks.

The rejection of claims 5, 12, and 20 is unsupported by the cited reference for at least the reasons given above in Argument A. Furthermore, Chambers fails to teach or suggest, "said second device is configured to calculate said delay value using one or more variables that correspond to one or more previous temporarily unavailable conditions" as recited in claim 5.

Applicant submits that Chambers teaches that the "delay input" corresponds to the known internal access latency of the target device. (Chambers, Column 5, Lines 44-51) Chambers teaches that the retry delay register, which is coupled to initiator PCI agent, receives and stores the "delay input" (i.e., the known latency of the target PCI agent) during a configuration access (e.g., system startup).

In support of his contention that Chambers teaches calculating the delay value using one or more variables that correspond to one or more previous temporarily unavailable conditions, the Examiner states, in the Final Office Action dated January 15, 2004, that this feature is taught in col. 6, lines 6-43 of Chambers. The Examiner specifies that Chamber teaches that the variable retry strategy is characterized as “n, 2, 2, 2...” (Chambers, column 6, lines 6-8); however, the variable “n”, which is referenced by the Examiner, corresponds to the known internal access latency of the target, which may vary from one target device to another. (Chambers, column 6, line 8, column 5, lines 44-51, and column 5, lines 59-60) Applicant respectfully submits that Chambers fails to teach or suggest that the delay value is calculated based on one or more variables that correspond to one or more previous temporarily unavailable conditions. For example, in one embodiment of Applicant’s invention, the delay value may vary depending on historical data from previous temporarily unavailable conditions. (Applicant’s Specification, Page 5, Lines 24-26)

In accordance, claim 5 is believed to patentably distinguish over the cited reference. Claims 12 and 20 recite features similar to those of claim 5 as discussed above and are thus also believed to patentably distinguish over the cited reference for at least the same reasons given above. Since the rejection is not supported by the teaching of the cited reference, Appellants respectfully request reversal of the Examiner’s rejection of claims 5, 12, and 20.

E. Claims 6, 13, and 18

The Examiner rejected claims 6, 13, and 18 under 35 U.S.C. § 102(e) as being anticipated by Chambers. Appellants respectfully traverse these rejections in light of the following remarks.

The rejection of claims 6, 13, and 18 is unsupported by the cited reference for at least the reasons given above in Argument A. Additionally, Chambers fails to teach or suggest, “said delay value corresponds to an encoded value” as recited in claim 6.

In support of his contention that Chambers teaches that the delay value is encoded, the Examiner states, in the Final Office Action dated January 15, 2004, that this feature is taught in col. 6, lines 6-43 of Chambers. Applicant respectfully submits that the variable “n”, which is referenced by the Examiner, corresponds to the known internal access latency of the target and it does not represent an encoded delay value. (Chambers, column 6, lines 6-8 and column 5, lines 44-51)

Therefore, Chambers fails to teach or suggest that the delay value corresponds to an encoded value. For example, in one embodiment of Applicant’s invention, “the delay value may be encoded to minimize the size and/or number of packets needed for the NAK. In one particular embodiment, the delay value can be encoded according to an exponential encoding in order to cover numerous orders of magnitude range.” (Applicant’s Specification, Page 5, Line 29 – Page 6, Line 3)

In accordance, claim 6 is believed to patentably distinguish over the cited reference. Claims 13 and 18 recite features similar to those of claim 6 as discussed above and are thus also believed to patentably distinguish over the cited reference for at least the same reasons given above. Since the rejection is not supported by the teaching of the cited reference, Appellants respectfully request reversal of the Examiner’s rejection of claims 6, 13, and 18.

F. Claims 7, 14, and 17

The Examiner rejected claims 7, 14, and 17 under 35 U.S.C. § 102(e) as being anticipated by Chambers. Appellants respectfully traverse these rejections in light of the following remarks.

The rejection of claims 7, 14, and 17 is unsupported by the cited reference for at least the reasons given above in Argument A. Furthermore, Chambers fails to teach or suggest, “a policy layer coupled to said first device and said second device, wherein said policy layer is configured to cause an error recovery mechanism to be initiated in response to detecting that a retry limit corresponding to said first request is exceeded, and wherein said error recovery mechanism is configured to perform an action according to said response” as recited in claim 7.

In support of his contention that Chambers teaches a policy layer which is configured to cause an error recovery mechanism to be initiated in response to detecting that a retry limit corresponding to the first request is exceeded, the Examiner states, in the Final Office Action dated January 15, 2004, and in the Advisory Action dated April 6, 2004, that this feature is taught in col. 7, lines 6-50 of Chambers. Applicant respectfully submits that col. 7, lines 6-50, of Chambers teaches a method for implementing the retry delay process of the invention. More specifically, Chambers teaches “The smart retry state machine 502 functions by intercepting the request input from the PCI agent 501 and subsequently issuing a delayed request to PCI bus 404 in its place.” (Chambers, column 7, lines 12-15)

Applicant further submits that Chambers teaches away from having a retry limit. More specifically, Chambers teaches that the variable retry strategy is characterized as “n, 2, 2, 2...”, which means that after initially waiting for the known internal access latency (i.e., the “delay input” or “n”) of the target device before the first retry access, the initiator device will retry every 2 PCI clock cycles. (Chambers, column 5, line 65 – column 6, line 2, and Figure 7, Block 708) Chambers teaches that the initiator retries every 2 PCI clock cycles and fails to teach about a retry limit.

Therefore, Chambers fails to teach or suggest an error recovery mechanism to be initiated in response to detecting that a retry limit corresponding to said first request

is exceeded. For example, in one embodiment of Applicant's invention, "If the sending device resends the request in excess of the retry limit, the policy layer may be configured to detect an error and may initiate an error recovery mechanism based on the type of NAK most recently received from the target device. The type of NAK may allow for different error recovery mechanisms based on different types of temporarily unavailable conditions at the target device." (Applicant's Specification, Page 6, Lines 16-20)


In accordance, claim 7 is believed to patentably distinguish over the cited reference. Claims 14 and 17 recite features similar to those of claim 7 as discussed above and are thus also believed to patentably distinguish over the cited reference for at least the same reasons given above. Since the rejection is not supported by the teaching of the cited reference, Appellants respectfully request reversal of the Examiner's rejection of claims 7, 14, and 17.

IX. CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejections of claims 1 – 20 were erroneous, and reversal of Examiner's decision is respectfully requested.

The Commissioner is authorized to charge the appeal brief fee of \$330.00 and any other fees that may be due to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5181-38400/BNK. This Appeal Brief is submitted in triplicate along with a return receipt postcard.

Respectfully submitted,



B. Noël Kivlin
Reg. No. 33,929
ATTORNEY FOR APPLICANT(S)

Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
(512) 853-8800
Date: 5-10-04

X. APPENDIX

The claims on appeal are as follows.

1. A computer system comprising:

a first device; and

a second device coupled to said first device;

wherein said first device is configured to convey a first request to said second device, wherein said second device is configured to receive said first request, wherein said second device is configured to detect a temporarily unavailable condition, wherein said second device is configured to convey a response to said first device corresponding to said first request, and wherein said response includes a delay value corresponding to said temporarily unavailable condition.

2. The computer system of claim 1, wherein said first device is configured to receive said response, and wherein said first device is configured to convey a second request to said second device at a time corresponding to said delay value.

3. The computer system of claim 1, wherein said second device is configured to generate said delay value according to a type of said temporarily unavailable condition.

4. The computer system of claim 1, wherein said delay value corresponds to a first value in response to said temporarily unavailable condition corresponding to a first type of condition and wherein said delay value corresponds to a second value in response to said temporarily unavailable condition corresponding to a second type of condition.

5. The computer system of claim 1, wherein said second device is configured to calculate said delay value using one or more variables that correspond to one or more previous temporarily unavailable conditions.
6. The computer system of claim 1, wherein said delay value corresponds to an encoded value.
7. The computer system of claim 1, further comprising:

a policy layer coupled to said first device and said second device, wherein said policy layer is configured to cause an error recovery mechanism to be initiated in response to detecting that a retry limit corresponding to said first request is exceeded, and wherein said error recovery mechanism is configured to perform an action according to said response.
8. A computer system comprising:

a communications medium;

a first device coupled to said communications medium; and

a second device coupled to said communications medium;

wherein said first device is configured to receive a response from said second device indicating that said second device is temporarily unavailable, wherein said response corresponds to a first request conveyed by said first device, wherein said response includes a delay value, and wherein said first device is configured to convey second request corresponding to said first request at a time corresponding to said delay value.

9. The computer system of claim 8, wherein said communications medium comprises a switching network.
10. The computer system of claim 8, wherein said communications medium comprises a shared bus.
11. The computer system of claim 8, wherein said communications medium comprises an arbitrated loop.
12. The computer system of claim 8, wherein said second device is configured to calculate said delay value using one or more variables that correspond to one or more previous temporarily unavailable conditions.
13. The computer system of claim 8, wherein said delay value corresponds to an encoded value.
14. The computer system of claim 8, further comprising:
a policy layer coupled to said communications medium, wherein said policy layer is configured to cause an error recovery mechanism to be initiated in response to detecting that a retry limit corresponding to said second request is exceeded, and wherein said error recovery mechanism is configured to perform an action according to said response.
15. A method comprising:
conveying a first request from a first device to a second device;
detecting a temporarily unavailable condition at said second device;

generating a delay value corresponding to said temporarily unavailable condition;
and

conveying a response corresponding to said first request to from said second
device to said first device, wherein said response includes said delay value.

16. The method of claim 15, further comprising:

conveying a second request from said first device to said second device at a time
corresponding to said delay value.

17. The method of claim 15, further comprising:

initiating an error recovery mechanism corresponding to said response in response
to determining that a retry limit corresponding to said first request has
been exceeded.

18. The method of claim 15, further comprising:

encoding said delay value prior to said conveying said response.

19. The method of claim 15, wherein said generating further comprises:

determining a type of said temporarily unavailable condition; and

generating said delay value according to said type of said temporarily unavailable
condition.

20. The method of claim 15, further comprising:

generating said delay value using one or more variables that correspond to one or more previous temporarily unavailable conditions.